

What is claimed is:

1 2. A method of restarting a node in a clustered computer system, wherein the
2 clustered computer system hosts a group including first and second members that
3 reside respectively on first and second nodes, the method comprising:

4 (a) in response to a clustering failure on the first node, notifying the
5 second member of the group the using the first member; and

6 (b) in response to the notification, initiating a restart of the first node
7 using the second member.

1 3. The method of claim 1, wherein the group comprises a cluster control
2 group that includes a member on each node participating in clustering in the clustered
3 computer system, and wherein the first and second members are each members of the
4 cluster control group.

1 4. The method of claim 1, wherein notifying the second member comprises
2 issuing a membership change request to the group using the first member.

1 5. The method of claim 3, wherein issuing the membership change request
2 includes indicating in association with the membership change request that the
3 membership change request is for the purpose of restarting the first node.

1 6. The method of claim 4, wherein indicating that the membership change
2 request is for the purpose of restarting the first node includes setting a reason field in
3 the membership change request to a restart value.

1 7. The method of claim 1, wherein initiating the restart includes issuing a start
2 node request to the group using the second member.

1 8. The method of claim 6, wherein issuing the start node request includes
2 indicating in association with the start node request that the start node request is for
3 the purpose of restarting the first node.

1 8. The method of claim 7, further comprising:
2 (a) detecting the clustering failure in the first node; and
3 (b) determining whether the clustering failure occurred during a restart
4 of the first node;
5 wherein notifying the second member of the group using the first member is
6 performed in response to detecting the clustering failure in the first node and
7 determining that the clustering failure did not occur during a restart of the first node.

1 9. The method of claim 8, further comprising signaling an error in response to
2 detecting the clustering failure in the first node if the clustering failure occurred
3 during a restart of the first node.

1 10. The method of claim 8, wherein determining whether the clustering failure
2 occurred during a restart of the first node includes determining whether the start node
3 request indicates that the start node request is for the purpose of restarting the first
4 node.

1 11. The method of claim 10, further comprising:
2 (a) counting protocols processed by the first node after a restart; and
3 (b) signaling an error in response to detecting the clustering failure in
4 the first node if the clustering failure occurred during a restart of the first node
5 and the number of protocols processed by the first node after the restart is less
6 than a predetermined threshold.

1 12. The method of claim 6, further comprising, in response to the clustering
2 failure on the first node, terminating clustering on the first node after notifying the
3 second member of the group using the first member.

1 13. The method of claim 1, further comprising, in response to the notification,
2 selecting the second member from a plurality of members in the group to initiate the
3 restart of the first node.

1 15. A method of restarting a node among a plurality of nodes in a clustered
2 computer system, wherein the clustered computer system hosts a cluster control group
3 including a plurality of cluster control members, each residing respectively on a
4 different node from the plurality of nodes, the method comprising:

5 (a) detecting a clustering failure on a first node among the plurality of
6 nodes;

7 (b) in response to detecting the clustering failure on the first node,
8 issuing a membership change request from the first node to the cluster control
9 member on each other node in the plurality of nodes, the membership change
10 request indicating that the membership change request is for the purpose of
11 restarting the first node;

12 (c) terminating clustering on the first node after issuing the
13 membership change request;

14 (d) in response to the membership change request, selecting a second
15 node from the plurality of nodes that is different from the first node;

16 (e) issuing a start node request using the selected second node, the
17 start node request indicating that the purpose of the start node request is for
18 restarting the first node; and

19 (f) in response to the start node request, initiating clustering on the
20 first node.

1 16. The method of claim 15, further comprising, in response to a second
2 clustering failure during initiation of clustering on the first node:

3 (a) determining from the start node request that initiated clustering on
4 the first node that the purpose of the start node request is for restarting the first
5 node; and

6 (b) in response to determining that the start node request is for
7 restarting the first node, signaling an error instead of initiating a second restart
8 of the first node.

1 17. An apparatus, comprising:
2 (a) a memory accessible by a node in a clustered computer system; and
3 (b) a program resident in the memory, the program configured to
4 initiate a restart of another node in the clustered computer system in response
5 to a notification from the other node of a clustering failure on the other node.

1 18. The apparatus of claim 17, wherein the program comprises a member of a
2 group hosted by the clustered computer system, the group including an additional
3 member residing on the other node.

1 19. The apparatus of claim 18, wherein the group comprises a cluster control
2 group that includes a member on each node participating in clustering in the clustered
3 computer system.

1 20. The apparatus of claim 18, wherein the program is configured to initiate
2 the restart by issuing a start node request to the group.

1 21. The apparatus of claim 20, wherein the start node request indicates that
2 the start node request is for the purpose of restarting the first node.

1 22. The apparatus of claim 17, wherein the program is configured to initiate
2 the restart of the other node responsive to a membership change request received from
3 the other node.

1 23. The apparatus of claim 19, wherein the program is configured to
2 determine whether the membership change request indicates that the membership
3 change request is for the purpose of restarting the first node.

1 24. A clustered computer system, comprising:

- 2 (a) first and second nodes coupled to one another over a network; and
3 (b) a group including first and second members, the first member
4 resident on the first node and the second member resident on the second node,
5 wherein the first member is configured to notify the second member in
6 response to a clustering failure on the first node, and wherein the second
7 member is configured to initiate a restart of the first node in response to the
8 notification.

1 25. The clustered computer system of claim 24, wherein the first member is
2 configured to detect the clustering failure in the first node, determine whether the
3 clustering failure occurred during a restart of the first node, and notify the second
4 member in response to detecting the clustering failure in the first node and
5 determining that the clustering failure did not occur during a restart of the first node.

1 26. The clustered computer system of claim 25, wherein the first member is
2 further configured to signal an error in response to detecting the clustering failure in
3 the first node if the clustering failure occurred during a restart of the first node.

1 27. The clustered computer system of claim 25, wherein the second member
2 is configured to initiate the restart by issuing a start node request to the group, and
3 wherein the first member is configured to determine whether the clustering failure
4 occurred during a restart of the first node by determining whether the start node
5 request indicates that the start node request is for the purpose of restarting the first
6 node.

1 28. The clustered computer system of claim 27, wherein the first member is
2 configured to signal an error in response to detecting the clustering failure in the first
3 node if the clustering failure occurred during a restart of the first node and a tracked
4 number of protocols processed by the first node after the restart is less than a
5 predetermined threshold.

1 29. The clustered computer system of claim 24, wherein the first member is
2 further configured to terminate clustering on the first node after notifying the second
3 member.

1 30. The clustered computer system of claim 24, further comprising a third
2 node coupled to the first and second nodes, wherein the group includes a third
3 member resident on the third node, wherein the first member is configured to notify
4 the third member in response to the clustering failure on the first node, and wherein
5 each of the second and third members is configured to locally select a single member
6 in the group to initiate the restart of the first node.

1 31. The clustered computer system of claim 30, wherein the second member
2 is configured to locally select the single member by determining whether the second
3 member is a lowest named member among the members of the group.

1 32. The clustered computer system of claim 24, wherein the group comprises
2 a cluster control group that includes a member on each node participating in clustering
3 in the clustered computer system, and wherein the first and second members are each
4 members of the cluster control group.

1 33. The clustered computer system of claim 24, wherein the first member is
2 configured to notify the second member by issuing a membership change request to
3 the group.

35. The program product of claim 34, wherein the signal bearing medium includes at least one of a recordable medium and a transmission medium.

36. A program product, comprising:

(a) first and second programs respectively configured to reside on first and second nodes in a clustered computer system, the first and second programs respectively operating as first and second members of a group, the first program configured to configured to notify the second program in response to a clustering failure on the first node, and the second program configured to initiate a restart of the first node in response to the notification; and

(b) at least one signal bearing medium bearing the first and second programs.

37. The program product of claim 36, wherein the first and second programs are borne on separate signal bearing media.

38. The program product of claim 36, wherein the first and second programs are borne on the same signal bearing medium.